REMARKS

Claim amendments

Applicant proposes amending the claims to recite *independent* first and second configuration databases. For reasons set forth below, Applicant believes this amendment renders the claims allowable over the cited art.

Support for the amendment can be found in the specification beginning on page 11, line 21.

Applicant also amends the claims to eliminate occurrence of phrases such as "at least a portion of" and "at least one." Since the claims use the open-ended transition "comprising," these phrases do not affect claim scope.

***Raz I* fails to teach second database with second configuration data**

In the preceding office actions, the Examiner has maintained that the second configuration data is disclosed in *Raz I* by the following passage on column 5, lines 25-55:

> "The importance of the mirroring feature to the present invention is that it enables the user to generate and maintain logically identical copies of portions of the database on two different disk devices within the data storage system.
>
> Notice however, that at any given point in time, the database that is being actively updated, and thus also the mirror copy, is not fully consistent. This is because at any given time some transactions have not yet completed and thus the data in the database associated with those transactions may still change during the course of completing those transactions. To achieve a fully consistent database, it is necessary to allow all of the transactions to be completed before being any new transactions are begun. When the pending transactions are completed, the database will be consistent. A consistent state can be achieved by gracefully shutting down the database, i.e., by preventing new transactions from being initiated and letting the pending transactions to run to completion. Typically, commercially available database systems allow the user to shut down the database gracefully to achieve full consistency. After the database has been shutdown, then the snapshot of the database (i.e., a copy of its contents at that time) will be consistent. The relevance of this will become apparent shortly.

BEST AVAILABLE COPY

> When data storage system 10 is configured to implement remote data mirroring, it generates and maintains a duplicate copy on the remote data storage system which is connected to the other end of high speed link 18. During remote data mirroring, the transfer of data over high speed link 18 operates in one of two modes, namely, a real-time mode and an asynchronous mode."[1]

Applicant has scrutinized the foregoing passage and found no hint of second configuration data. The passage does state that during mirroring, there are two databases that are briefly inconsistent. Applicant suspects that the Examiner infers the existence of second configuration data because:

1. The first database contains configuration data;
2. There exists a second database that is a mirrored copy of the first database; and
3. Therefore, the second database must contain configuration data that is a copy of that found on the first data.

Applicant disagrees with the first proposition, namely that the first database contains configuration data. Nevertheless, in an effort to expedite prosecution, Applicant amends the claims to recite first and second configuration data that are independent of each other.

The independence of the first and second configuration data in the amended claims means that the hosts having access to the primary system can be different from those that have access to the backup system. This feature is neither disclosed nor suggested by *Raz I*.

Raz I fails to teach first database with first configuration data

In support of the proposition that there exists a first database that includes first configuration data, the Examiner has consistently drawn attention to column 7, lines 21-60, which reads:

> "In this example, we assume that the decision-support functions are performed by using the same data storage system through which the database updating functions are performed. In other words, the portion of system shown in FIG. 1 which is relevant to

[1] *Raz*, et al., U.S. Patent No. 5,852,715, ("*Raz I*") col. 5, lines 25-56.

BEST AVAILABLE COPY

this example is data storage system **10** with the plurality of host processors that are connected to it. We further assume that host processor **14(1)** is performing updating functions to a database that is stored in a volume **80** within the data storage system. Data storage system **10** is configured to permit the other host processors **14(2)-14(n)** to share volume **80** with host processor **14(1)**. Of course, it could also be the case that the database is spread across multiple volumes in data storage system **10** in which case, data storage system **10** would be configured to permit all of the other host processors to share all of those volumes with host processor **14(1)**. But for simplicity we will assume that the database is contained within a single volume.

In accordance with the first example, host processor **14(1)** is permitted to both write data to and read data from the shared volume **80**. In contrast, the other host processors **14(2)-14(n)**, by a restriction that may be enforced either at the host processor level or within the data storage system itself, are only permitted to read data from shared volume **80**. Thus, for example, host processor **14(1)** might be performing regular database work, e.g. online transaction processing or excepting update batches; while the remaining host processors **14(2)-14(n)** are performing decision-support functions, for which only read capability is required and used.

Since only one host processor is writing to the database, no lock manager is needed to coordinate the accesses to the database by the other host processors. More specifically, though there may need to be intrahost locks for read/write accesses coming from the one host processor, there need not be interhost locks for the other host processors. Thus, reading of the data within the database for purposes of decision-support can take place concurrently with the updating of the database bvy the one host processor. That is, all of the host processors can operate concurrently, assuming that the above-described restriction on only allowing host processor **14(1)** to perform updating is honored. Stated differently, the read accesses by the host processors that are performing decision-support can be done concurrently with the writing of data by host processor **14(1)**."[2]

Despite having repeatedly scrutinized this passage, Applicant is unable to identify where precisely there is a teaching of a first database that *includes* the first configuration data.

As best understood, the Examiner's reasoning for inferring the existence of the claimed first database with first configuration data on the basis of the foregoing passage is as follows:

1. The foregoing passage refers to the existence of a first database.

2. It also discloses that a host can access only a subset of volumes.

---

[2] *Raz I*, col. 7, lines 21 to 67.

BEST AVAILABLE COPY

3. Therefore, someplace there must exist configuration data that specifies which hosts can access which volumes.

4. Therefore, this configuration data must reside in the first database.

Applicant agrees that conclusion 3 follows from observations 1 and 2. However, Applicant questions the conclusion of step 4. The foregoing passage does not state where the configuration data might be kept. The fact that configuration data exists *someplace* in the system does not mean it exists in the first database.

Where then is this configuration data kept? *Raz I* suggests an answer at column 4, lines 49-54, which states:

> "Data storage system **10** can be configured into multiple logical volumes. Typically, a volume corresponds to a single disk device. A service console **50** within data storage system **10** enables the user to configure the data storage, i.e., to define the logical volumes and to specify which logical volumes are accessible through which host connections **20**."[3]

The foregoing passage suggests that the configuration data can be found on the service console **50**. Referring to FIG. 2, we see that the service console **50** is distinctly *separate* from the volumes **22** on which the database resides. Therefore, the database cannot be said to *include* the configuration data.

***Raz II* fails to teach independent second configuration data**

Raz II teaches a system in which each host maintains a copy of the ownership table **38**. This is referred to in column 4, lines 52-58, which states:

> "All host processors have access to both the database table and the ownership table in the locally accessible memory. Alternatively, each host processor can keep its own *copies* of each of these tables. The approach that is selected in any given system is an implementation detail." [*emphasis supplied*][4]

---

[3] *Raz I*, col. 4, lines 49-54.
[4] *Raz*, et al., U.S. Patent No. 5,860,137, ("*Raz II*") col. 4, lines 52-58.

BEST AVAILABLE COPY

PAGE 9/12 * RCVD AT 6/21/2004 2:36:26 PM [Eastern Daylight Time] * SVR:USPTO-EFXRF-1/2 * DNIS:8729306 * CSID:617 542 8906 * DURATION (mm-ss):03-56

Raz II therefore teaches a system in which there are multiple copies of the *same* ownership table. The ownership tables are therefore *not independent* of each other. Therefore, *Raz II* does not teach *independent* first and second configuration data as recited in the amended claims.

There is no reason to suppose one would modify the *Raz II* system so that ownership tables could be independent of each other. This would surely invite chaos, as hosts attempt to access volumes that other hosts expect to have exclusive access to.

***Raz II* fails to teach backup system**

The independent claims require a backup system having a backup storage device. The backup storage device stores a backup copy of data stored on the primary storage device.

There is no doubt that *Raz II* discloses a backup management processor. However, the backup management processor does *not* store a copy of data stored on in the data storage system 14. Therefore, the backup management processor cannot correspond to the backup system recited in the claims.

It is true that the backup management processor maintains a copy of the ownership table 38 and the database table 36. However, these tables are not stored on the data storage system 14. They are both stored in a memory that is local to the management processor. This is readily apparent both from FIG. 2, and from column 4, lines 49-53:

> "With this information, the managing host processor generates in a locally accessible memory 31 a database table 36 which maps the database records to the logical volumes 20 and an ownership table 38 which assigns ownership of the logical volumes to the host processors 12." [*emphasis supplied*][5]

Consequently, the backup management processor cannot possibly be a backup system having a backup storage device that stores a backup of data stored on the primary storage device. The only data that the backup management processor stores is a copy of data that is in the local memory of the management processor, and *not* on the primary storage device.

[5] *Raz II*, col. 4, lines 49-53.

BEST AVAILABLE COPY